

SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

Scalable Video Encoder/Decoder with Drift Control

Cross Reference to Related Applications

This application claims priority to United States Provisional Application "MANAGING DRIFT IN DCT-BASED SCALABLE VIDEO CODING," Serial No. 60/252,525, filed on November 22, 2000, the contents of which are incorporated by reference herein.

Background of Invention

[0001] The present invention relates to video compression technology and, more particularly, to scalable video coders.

[0002] Compressed video, which uses predictive coding algorithms and variable length coding, is sensitive to network impairments since these can cause error propagation. A single bit error or erasure can cause substantial degradation if no action is taken to stop or limit the extent of error propagation. Motion compensation allows the error to propagate both temporally and spatially. Because of this, there has been extensive effort in the video community to design new techniques that limit the extent of error propagation. However, almost all attempts to limit error propagation decrease the coding efficiency, some dramatically so. To ensure the best operation of the video coder in an errorprone channel, the balance between resilience and efficiency must be managed carefully.

[0003] Scalable coding algorithms create a partitioning of the compressed bitstream into more and less important parts. This allows a natural combination with different mechanisms to prioritize network transport, for example, marking less important parts for early discard, applying unequal error protection, or facilitating rate matching between encoder and network. When used in conjunction with such techniques,

scalable video can be very resilient to networkintroduced errors.

[0004] The propagation of enhancement-layer errors into the base-layer reconstruction is referred to herein as "drift" (this is in contrast to the more general term "error propagation" which is used herein to include the result of partial reception of the more important base-layer information). Early scalable video coders (like MPEG2 SNR scalability (SNRS)) allowed drift by using lowpriority enhancement-layer information to predict the highpriority base-layer. However, in recent years, the development of scalable video encoders (like H.263 SNRS and spatial scalability (SS)) has focused on eliminating drift. In these algorithms, the base-layer is predicted only from the base-layer information. This strategy has been taken one step further in the development of MPEG4 Fine Granularity Scalability (FGS), in which the enhancement-layer information is also predicted only from base-layer information.

[0005] However, while recent scalable video coding algorithms are becoming more efficient at compressing the video, they lose compression efficiency because they ignore all enhancement-layer information when predicting the base-layer. In particular, recent experiments show that with MPEG2 SS, MPEG4 and H.263 scalability modes all suffer from 0.5 – 1.5 dB losses for every layer. FGS has particularly poor compression inefficiency because of its restricted prediction strategy.

Summary of Invention

[0006] Despite the predominance of arguments in the prior art for designing video compression systems that eliminate drift, the present invention is based on the observation that if one can effectively manage error propagation in both a one-layer encoder and a two-layer encoder that does not allow the introduction of base-layer drift, it is possible to design an encoder that does allow the introduction of drift into the base-layer.

[0007] It is an object of the invention to provide means to introduce drift incrementally.

[0008] It is an object of the invention to provide a way for the encoder to measure the drift being potentially introduced.

[0009] It is an object of the invention to provide encoding options that can allow drift,

while simultaneously keeping the amount of drift under control.

- [0010] It is an object of the invention to provide means to drastically reduce or eliminate drift without the need for a full I - frame.
- [0011] It is an object of the invention to provide a system-level optimization, designed to maximize expected quality across all expected receivers.
- [0012] The present invention is directed to structures for scalability that use previous enhancement-layer information in a video stream to predict the current base-layer, while simultaneously managing the resulting possibility of drift. These structures allow better compression efficiency, while introducing only limited impairments in the quality of the reconstruction.
- [0013] These and other advantages of the invention will be apparent to those of ordinary skill in the art by reference to the following detailed description and the accompanying drawings.

Brief Description of Drawings

- [0014] FIG. 1 is a block diagram of a two-loop decoder with drift control, embodying aspects of the invention.
- [0015] FIG. 2 is a block diagram of a two-loop encoder with drift control, embodying aspects of the invention.
- [0016] FIG. 3 is a flowchart illustrating the operation of the drift compensation component of the decoder, in accordance with an embodiment of the invention.
- [0017] FIG. 4 is a flowchart illustrating the operation of the encoder in selecting effective encoding parameters, in accordance with an embodiment of the invention.

Detailed Description

- [0018] FIG. 1 sets forth a simplified block diagram of a scalable DCT - based motion-compensated video decoder with drift control, illustrating an embodiment of the present invention. The decoder takes three levels of input, as shown in FIG. 1. The base bits 110, with bit rate

R_{nc}

are assumed to be always available. The enhancement bits are divided into a first part 111 and a second part 112. The first part of the enhancement bits 111, with bitrate

$R_{nf} - R_{nc}$

may not be received by the decoder, but if received, are used to predict the next frame. The second part of the enhancement bits 112, with bit rate

$R_{na} - R_{nf}$

may not be received, and is never used to predict the next frame. The base layer bits are assumed to include control information 151, i.e. motion vectors for motion compensation and macroblock type information, the use of which is described in further detail herein.

PCT/GB2013/05303

[0019] The bit stream inputs of FIG. 1, first, are used at 120, 121, 122 to extract and decode the variable length coded (VLC) words to obtain locations and quantizer values of the DCT coefficients for each block. The non-zero DCT coefficients are reconstructed at 125 from the base layer. At the base layer, the DCT coefficients are coarsely quantized. The enhancement layer information encodes the difference between the non-quantized DCT coefficients and the quantized coefficients from the base layer with finer quantization stepsize. These differences, if available, are reconstructed at 126 and 127 and combined at adders 161 and 162 with the base layer reconstruction to obtain the DCT coefficients. Inverse discrete cosine transform (DCT) 132 is used to obtain the quantized block pixel values. Motion compensation is performed and provided at adder 164. By processing the bit streams of both the base layer and enhancement layer, the highest quality video signal 100 may be reconstructed.

[0020] Both the decoder shown in FIG. 1 and the encoder shown in FIG. 2 maintain two frame memories: referred to by the inventors as a "coarse" frame memory and a "fine" frame memory. The coarse and fine frame memories are shown in FIG. 1 as 140 and 141, with corresponding motion compensation 145, 146 and inverse DCT 130, 131. The "coarse" frame memory 140 depends only on the base bits and never drifts. The "fine" frame memory 141 is updated by first combining both motion compensated frame memories, and then applying the base bits and the first part of the

enhancement bits. The fine memory drifts when some of these enhancement bits are lost.

[0021] The motion-compensated predictions from the coarse and fine memories 140 and 141 are referred to as

\hat{p}_{nc}

and

\hat{p}_{nf}

for a macroblock n and are labelled as 155 and 156 respectively in FIG. 1. For simplicity, the same motion vectors can be used for both predictions. For each macroblock, the drift compensation component 150 in FIG. 1 combines the coarse and fine predictions according to a macroblock type information. As mentioned above, the macroblock type information is included in the control information 151 of the base layer bits. The macroblock type information is utilized to control and manage the introduction of drift into the base layer.

[0022] FIG. 3 is a flowchart illustrating the operation of the drift compensator 150. The coarse and fine predictions are input at step 301. At step 302, the macroblock control information is read. For simplicity, only three drift management options are shown in FIG. 3. The first option, at steps 303–304, eliminates drift by taking the coarse prediction

\hat{p}_{nc}

only (as in FGS). The second option, at steps 305–306, allows drift by taking the fine prediction

\hat{p}_{nf}

only (as in MPEG2 SNRS). The third option, at steps 307–308, reduces -- but does not eliminate -- drift by averaging both predictions

$(\hat{p}_{nc} + \hat{p}_{nf})/2$

Although only three options are shown in FIG. 3, introducing new combinations would naturally extend the present scheme.

[0023] FIG. 2 sets forth a simplified block diagram of a scalable DCT video encoder with drift control, corresponding to the decoder in FIG. 1. A video sequence is input at 200.

The DCT 230 is applied to each block and the DCT coefficients are quantized at 229. As mentioned above, the DCT coefficients are coarsely quantized at the base layer using quantizer Q_c . The enhancement layer information encodes the difference between the non-quantized DCT coefficients and the quantized coefficients from the base layer with finer quantization stepsizes. The base layer bits 210 are created by coding the coarse quantizer values using variable length code (VLC) tables 220. The baselayer VLC as shown in FIG. 2 relies on arithmetic bitplane coding, but could also be implemented, without limitation, using the conventional Huffman method. The results of the other quantizers Q_f and Q_a are also bit-plane coded at 221 and 222 to produce the first and second enhancements bits, 211 and 212 respectively. The first frame of the video sequence (an I - frame) is encoded without reference to any past or future frames. Motion compensation is utilized at adder 261 with regard to subsequent P - frames. The encoder tracks a coarse and fine frame memory 240 and 241, corresponding to the coarse and fine frame memories 140 and 141 of the decoder. The frames stored in the coarse and fine frame memories 240 and 241 are reconstructed using inverse DCT 230, 231 and inverse quantizers 225, 226. Corresponding motion compensation 245 and 246 are utilized to obtain the coarse and fine predictions 255 and 256, which are utilized by the drift compensator 250, as described above, to compute a combination based on control information 251. A drift control/motion estimator 270 is utilized by the encoder to provide control information 251 to the quantizers, to the motion compensation units, to the drift compensator, and for encoding in the base layer. The motion estimator 270 chooses motion vectors, as is well understood in the art, which are included in the control information encoded in the base layer. The drift control 270, using the video input 200 and the frame memories 240, 241, chooses a particular combination mode for the drift compensator 250 and adjusts quantization in an advantageous manner, as further described herein.

[0024]

To minimize the influence of drift in general, it is advantageous to use an embedded coder to compress each individual frame. This allows more significant enhancement layer bitplanes to be received and decoded even if the network does not have sufficient bandwidth to send the entire enhancement layer. A scalable video bitstream must indicate the relative importance of bits or groups of bits within the bitstream. This information is necessary to facilitate transport mechanisms that

respond to the observed channel characteristics by discarding the least important bits or by applying error correction to the more important bits. In FIG. 2, a scalable bitstream is produced by a binary adaptive Z - Coder, see L. Bottou et al., "The Z - Coder Adaptive Binary Coder," Proc. IEEE Data Compression Conference 1998, which is incorporated by reference herein, and prioritization is imposed by the order in which information is encoded during a "coding run".

[0025] The Z - Coder coder represents probability distributions using context variables. Internal state and context variables are initialized at the beginning of each coding run and are updated after coding each binary decision. If a piece of the bitstream is missing, the decoder cannot perform identical updates and must stop. Decoding can resume at the beginning of the next coding run. Hence, more important information should be encoded early in the coding run to ensure it is received prior to any information loss.

[0026] Shorter coding runs increase robustness to random losses in the channel. However, where prioritized transport is assumed, very short coding runs will not be useful in the system. Longer coding runs typically improve the coding efficiency because the contexts do not need to be relearned at the beginning of every coding run. A typical single frame requires 1.072 or 0.513 bits per pixel when quantized with $Q = 4$ or $Q = 8$ respectively. If the same frame is encoded 256 times without resetting the adaptation contexts after encoding each copy, the resulting bitrates are 1.069 and 0.509 bits per pixel respectively. This suggests that no more than 1% of the bandwidth could be saved by coding several frames per coding run. Further, latency becomes problematic if there are several frames per coding run. The inventors have selected for their experiments a scalable coder which performs one coding run per frame.

[0027] Within each coding run, binary decisions are encoded by decreasing order of importance. We start with the macroblock type, quantization and motion information, followed by the DCT coefficient information for the base layer and the various enhancement layers. The DCT coefficients are encoded as a sequence of binary decisions ordered according to their importance. The coefficients in each 8x8 block are first divided by the finest quantizer Q_a , resulting in an integer quotient with

absolute value that can always be represented using twelve bits. Each elementary bitplane coding operation then processes one particular bitplane of one particular 8x8 block of DCT coefficients. Each elementary coding operation belongs either to the base layer or to one of the two refinement layers, according to the bitplane number and to the values of the $Q_c = Q_a$ and $Q_f = Q_a$ ratios for the corresponding macroblock. First, the base layer is encoded by performing all the baselayer elementary coding operations starting with the most significant bitplanes, and proceeding towards the least significant bitplane. Then each refinement layer is encoded by similarly performing all the elementary coding operations for that layer in decreasing order of bitplane significance. The Z - Coder internal state is flushed after each layer in order to clearly delimit the bitstream for each layer.

[0028] To perform an elementary coding operation, the coder loops over the 64 coefficients and codes whether the binary representation of its absolute value contains a 0 or a 1 for the current bitplane. The sign bit is coded just after coding the first 1 of each coefficient. Such coefficients are named significant. Previous bit plane coders take advantage of the wavelet transform structure by coding decisions addressing the significance of entire coefficient groups. The same result is achieved for zig - zig ordered DCT coefficients by coding a stopping decision after each significant coefficient. A positive stopping decision indicates that none of the remaining coefficients will be significant after processing this bitplane. When it is known that some remaining coefficients are already significant, there is no need to code the stopping decision.

[0029] This encoding process produces a high bandwidth bitstream. Lower bandwidth channels are accommodated by simply truncating a fixed proportion of the second refinement layer. Even lower bandwidth channels are accommodated by eliminating the second refinement layer altogether and truncating a fixed proportion of the first refinement layer. Extremely low bandwidth channels might be dealt with by eliminating both refinement layers and truncating the base layer. Here, the truncation process is facilitated by dividing the bitstream into "chunks", with one chunk per layer per frame. Each chunk is preceded by a one- to three-byte length indicator, indicating how many bytes are in the chunk. There are three chunks per frame, corresponding to the base layer, the first part of the enhancement layer, and the second part of the

enhancement layer.

[0030] The encoder in FIG. 2 tracks both frame memories under the assumption that all bits are received by the decoder. The encoder makes several decisions that affect the decoder drift in the fine memory. The first decision is the selection of a combination mode for the drift compensation. The second decision involves the number of bit-planes that might be used in the prediction loop; this is accomplished by adjusting the quantization Q_f relative to the final quantization Q_a . Although not shown in FIG. 2, a third technique could be to apply a filter to the prediction from the coarse loop in order to smooth the discontinuities associated with prediction blocks that straddle macroblock boundaries.

[0031] Different images have different tradeoffs between efficiency and resilience as a function of these drift control decisions. The encoder must make these decisions and send this information to the decoder. The encoder makes these choices on a macroblock basis with the goal of optimizing the total system performance.

[0032] The traditional (often implicit) optimization when designing a scalable coder is to minimize the maximum possible distortion at the decoder, subject to the constraint that the channel rate R is

$$\bar{R}_c \leq R \leq \bar{R}_a$$

Typically, both

$$\bar{R}_c$$

and

$$\bar{R}_a$$

are known, although neither the instantaneous channel rate nor the average channel rate in some time interval is known. This maximum distortion is achieved for the minimum rate

$$\bar{R}_c$$

Thus, optimizing using this criterion leads to a prediction process that does not tolerate any drift in the base layer. However, this also results in larger than necessary distortions at rates near

$$\bar{R}_a$$

Following is a discussion of some alternate criteria for optimization, to achieve better compression at high rates without sacrificing too much quality at lower rates.

[0033] One optimization criterion is to minimize the distortion at the highest rate, subject to constraint that the drift at the lowest rate is kept below some value. This can be expressed as

$$\begin{aligned} & \min\{D_a\} \\ & \text{subject to } D_c \leq \bar{D}_c \\ & \text{and the rate constraints} \\ & R_c \leq \bar{R}_c \quad \text{and} \quad R_a \leq \bar{R}_a \end{aligned}$$

Here, R_c and R_a are the rates associated with the base bits, and all the bits, respectively, and D_c and D_a are the distortions of the associated reconstructions.

[0034] An alternate optimization criterion is to minimize the distortion averaged across all possible channel rates, subject to the rate constraints shown above. Determining the average distortion, however, requires knowledge of the probability distribution of the rates. This information is typically not available. However, a simple approximation is to minimize a weighted distortion

$$w_c D_c + (1 - w_c) D_a$$

subject to the rate constraints set forth above. The choice of the weights is influenced by the application

[0035] These two constrained optimizations can be solved by converting them to unconstrained Lagrangian optimizations. The unconstrained optimization problems will have two or three Lagrangian parameters, and can be solved using conventional techniques. See, e.g., S. Wenger et al., "Using RFC2429 and H.263+ at low to medium bit-rates for low-latency applications," Packet Video Workshop '99; R. Zhang et al., "Video coding with optimal Inter/Intra - mode switching for packet loss resilience," IEEE Journal on Selected Areas in Communications, vol. 18, no. 6, pp. 966–976, June 2000; R. Zhang et al., "Switched error concealment and robust coding decisions in scalable video coding," IEEE International Conference on Image Processing, Oct. 2000.

[0036] Even using a simple optimization approach demonstrates the benefits of the present invention. For example, consider the following simple heuristic where the quantizers are fixed. With the desire to keep the amount of drift limited, each

macroblock is coded based on the sum of the absolute error of the three possible predictions: S_{nc} , S_{nf} and $S_{n(c+f)}$. If the drift for this macroblock will be larger than some threshold,

$$S_{nc} \geq T_{drift}$$

the macroblock is coded using the prediction

$$(\hat{p}_{nc} + \hat{p}_{nf})/2$$

Otherwise, the prediction that has the minimum sum of the absolute error is chosen.

An encoder using this simple heuristic described above can be shown to outperform a range of other coders across a range of channel rates.

[0037] A more detailed optimization approach based on locally adaptive selection of quantizer and the prediction is herein described. Rather than fixing the quantizers and using a simple heuristic to choose which prediction should be used for each macroblock, both the quantizer Q_f and prediction are changed with the goal of minimizing

$$\sum_t R_f(t)$$

subject to

$$\sum_t [(1-\gamma) \log D_f(t) + \gamma \log D_c(t)] \leq \sum_t \log D^*$$

where $R_f(t)$ and $D_f(t)$ are the rate and distortion produced when using $Q_f(t)$, and $D_c(t)$ is the distortion with only the coarse information, each at time step t . To simplify the notation, it is helpful to suppress the fact that the summation applies not only over all time steps, but also over all macroblocks in each frame.

[0038] The value of γ depends on the application, and should be chosen to create a network-aware encoding. If the channel rate is expected to be generally close to R_{\min} , then the parameter should be set close to 1. Then, the encoder will reduce the drift at the lower rates at the expense of the compression efficiency at higher rates. On the other hand, if the channel rate is expected to be generally greater than R_{\min} with occasional dips, then the parameter should be set close to zero.

[0039]

Applying Lagrangian optimization, we choose Q_f and the prediction method p such that

$$\min_{\{p, Q_f\} \forall t} \sum_t [\mu R_f(t) + (1 - \gamma) \log D_f(t) + \gamma \log D_c(t)]$$

Very simple models for rate and distortion may be used:

$$R = \frac{1}{2} \log(\sigma^2 / D), \quad D_f = \frac{Q_f^2}{12}, \quad \text{and} \quad D_c = \frac{Q_c^2}{12} + M_p,$$

where M_p is a measure of the amount of mismatch between a given prediction p and the coarse prediction

$$\hat{p}_{nc}$$

to account for the contribution of drift to D_c . Using these, the optimization equation becomes:

$$\min_{\{p, Q_f\} \forall t} \sum_t \left[\mu \log \sigma_p^2(t) + (1 - \mu) \log \frac{Q_f(t)^2}{12} + \gamma \log \left(\frac{Q_c(t)^2}{Q_f(t)^2} \right) + \gamma \log \left(1 + \frac{12M_p(t)}{Q_c(t)^2} \right) \right]$$

where p indicates the prediction being evaluated, and the constant multiplier has been absorbed into μ .

[0040] In general, the first and fourth terms in the above equation depend on past predictions and quantizers, as well as on the prediction in the current time step. For example, consider the selection of the best predictor for a particular macroblock in frame t . This macroblock depends on up to four MBs in frame $t - 1$. If a finer quantizer Q_f had been used for those four MBs in frame $t - 1$, then the prediction

$$\hat{p}_f$$

in frame t would have had a smaller

$$\sigma_p^2$$

but larger M_p . Because of this dependency across time, joint optimization is required; however, a fully joint optimization requires a prohibitive search space.

[0041]

Here, it is possible to ignore some of the dependencies to simplify the optimization procedure. Because the dependence on the previous prediction is generally weaker than the dependence on the previous quantizer, it is possible to first consider jointly only the quantizer selection at time $t - 1$ with the prediction selection at time t . However, this is generally still too complicated because one MB in frame t depends on four different values of Q_f in frame $t - 1$. This is still too large a state

space to reasonably consider.

[0042] Hence, it is advantageous to consider first the choice of the best quantizer for each MB in frame $t - 1$ assuming the neighboring quantizers are identical, and consider second the choice of the best prediction using the selected quantizers.

[0043] Consider the second step first. Assuming the quantizers Q_f and Q_c are fixed for all time and that the current predictor depends only weakly on the previous predictors, the above optimization equation reduces to

$$\min_{p(t)} \mu \log \sigma_p^2(t) + \gamma \log \left(1 + \frac{12M_p(t)}{Q_c(t)^2} \right)$$

for frame t . Thus, to choose the best predictor given fixed quantizers, one can simply evaluate this equation for each of the three predictors

$$\hat{p}_{nc}$$

$$\hat{p}_{nf}$$

$$(\hat{p}_{nc} + \hat{p}_{nf})/2$$

and choose the minimizing predictor. Note that in the case when $\gamma = 0$, this simply chooses the predictor which has the minimum prediction error. If $\gamma > 0$, the impact of drift is also considered.

[0044] The first step, to choose the best quantizer for a MB in frame $t - 1$ assuming the prediction strategies are already fixed, is more complicated. One needs to consider the first and fourth terms of the optimization equation for frame t , and the second and third terms of the optimization equation for frame $t - 1$. One can ignore the second and third terms for frame t because the quantizer Q_f at time t is unknown, and can be assumed to be constant. The first and fourth terms for frame $t - 1$ can be ignored because the predictor for $t - 1$ is already fixed. Thus, to choose the quantizer $Q_f(t - 1)$, one must minimize:

$$(1 - \mu) \log Q_f(t-1)^2 / 12 + \gamma \log \left(\frac{Q_c(t-1)^2}{Q_f(t-1)^2} \right)$$

$$+ \mu \log \sigma_p^2(t) + \gamma \log \left(1 + \frac{12M_p(t)}{Q_c(t)^2} \right)$$

[0045] For this first step, one must consider jointly the effect of the quantizer $Q_f(t-1)$ and the predictor $p(t)$. A particular MB in frame $t-1$ may be used to predict multiple MBs in frame t . Therefore, all affected MBs in frame t are first determined via reverse motion compensation. For every possible quantizer $Q_f(t-1)$, it is assumed that the surrounding MBs use the same quantizer, and a determination is made of the best prediction of those affected MBs, as in the above equation for predictors. Then, the $Q_f(t-1)$ is chosen that minimizes the weighted cost of the affected MBs. To compute the appropriate weight, one can use the number of pixels in each MB in frame t that are impacted by the current MB in frame $t-1$.

[0046] FIG. 4 sets forth a flowchart, illustrating an implementation of the above algorithm. Steps 401 and 402 are preprocessing steps, in which motion estimates are made and optimal predictors are chosen based on the predictor optimization equation above. From steps 403 to 413, each macroblock in frame t is considered. For each affected macroblock in a future frame $t+1$ and for each quantizer $Q = m Q_a$ (where $m = 1, 2, 4, 8$), the predictor which minimizes the predictor optimization equation is determined which is used to compute the incremental cost for this affected macroblock, at steps 404 to 410. At step 411, the quantizer is chosen which has the minimum cumulative weighted cost. At step 412, the current macroblock is coded using the chosen quantizer and predictor.

[0047] In a current implementation of the drift-controlled coder, the inventors have set $Q_a = 4$ and $Q_c = 8Q_a$, with the choice of Q_f and the prediction as described above. Such an implementation can be shown to significantly outperform alternatives across a range of channel rates without incurring significant drift for low channel rates.

[0048] The embodiment of the present invention described above enables effective drift management by incorporating the following features:

- First, it advantageously provides a means to introduce drift incrementally.

This is straightforward to achieve by bitplane encoding or by creating an embedded bitstream. In order to be effective, this should be used in conjunction with a mechanism in the transport for providing more reliable delivery of the more important bitplanes to the receiver. See, e.g., P.A. Chou, et al., "FEC and pseudo-ARQ for receiver-driven layered multicast of audio and

video," Data Compression Conference, pp. 440–449, Mar. 2000; R. Rejaie, et al., "Quality adaptation for congestion controlled video playback over the Internet", Proceedings of ACM SIGCOMM '99, Cambridge, MA., Sept. 1999; which are incorporated by reference herein.

- Second, it advantageously provides a way for the encoder to measure the drift being potentially introduced, so that the encoder knows when drift is becoming significant.
- Third, it advantageously provides encoding options that can allow drift (i.e., allow errors in the enhancement layer to propagate into the base layer), while simultaneously keeping the amount of drift under control.
- Fourth, it advantageously provides a means to drastically reduce or eliminate drift without the need for a full Iframe.
- Fifth, it advantageously provides a systemlevel optimization, designed to maximize expected quality across all expected receivers. Inherent to this optimization, there should be some (possibly inaccurate) knowledge on the part of the encoder as to how many errors the channel will introduce, and how those errors will be introduced (gradually bitplane by bitplane, or suddenly when an entire packet of highpriority data is lost).

It should be noted that partial management of drift is still possible with different subsets of the above features.

[0049] The foregoing is to be understood as being in every respect illustrative and exemplary, but not restrictive, and the scope of the invention disclosed herein is not to be determined from the Detailed Description, but rather from the claims as interpreted according to the full breadth permitted by the patent laws. Embodiments within the scope of the present invention also include device readable media and computer readable media having executable program instructions or data fields stored thereon. Such computer readable media can be any available media which can be accessed by a general purpose or special purpose computing device.

[0050] It is to be understood that the embodiments shown and described herein are only illustrative of the principles of the present invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention. For example, and without limitation, while the above discussion

focuses on hybrid block-based DCT approaches to layering, the present invention can be readily extended to wavelet video coders that use some form of motion-compensated prediction. Furthermore, although the invention is described herein with respect to a decoder and encoder that has three levels of input/output and has two frame memories, etc., the present invention is not so limited and may be readily extended to other numbers of inputs, frame memories, and other components by one of ordinary skill in the art. Also, the above discussion focuses on ways to manage drift within the predictive framework of P - frames. Structures such as B - frames or P ' - frames (which are similar to B - frames without forward prediction, and which are enabled by Reference Picture Selection (RPS) mode of annex N in H.263+) are not explicitly considered, even though these structures naturally reduce drift by having fewer predictions made from partially correct data. B - and P ' - frames can easily be incorporated into the present system level structure, and indeed, a P ' - frame is one way to limit the temporal extent of error propagation without an I - frame, even for a one-layer encoder.